

# HP StorageWorks Modular Smart Array 500 G2

## technology

technology brief



|  |    |
|--|----|
| Abstract.....                                  | 3  |
| Introduction.....                              | 3  |
| Technology overview .....                      | 4  |
| MSA500 G2 controller.....                      | 4  |
| Redundant controller operation .....           | 5  |
| Recovery ROM and ROM cloning .....             | 6  |
| I/O module.....                                | 6  |
| Power supplies and cooling.....                | 6  |
| Host bus adapters .....                        | 6  |
| Configuration flexibility .....                | 7  |
| MSA500 G2 High Availability Kit.....           | 7  |
| 2-node multipath configuration.....            | 8  |
| 4-node shared storage configuration .....      | 8  |
| 4-node mixed configuration .....               | 9  |
| ProLiant DL380 G4 Packaged Cluster system..... | 9  |
| Selective Storage Presentation .....           | 10 |
| Availability .....                             | 10 |
| Advanced Data Guarding.....                    | 10 |
| Distributed data guarding .....                | 10 |
| Drive mirroring .....                          | 10 |
| Online spares.....                             | 10 |
| Management technologies.....                   | 11 |
| Hardware and software compatibility .....      | 12 |
| Performance enhancement .....                  | 12 |
| Data migration .....                           | 13 |
| DAS-to-SAN conversion .....                    | 13 |
| Conclusion.....                                | 14 |

|                           |    |
|---------------------------|----|
| For more information..... | 14 |
| Call to action .....      | 14 |

## Abstract

The HP Modular Smart Array 500 G2 (MSA500 G2) system integrates many high availability features into a familiar, easy-to-deploy and easy-to-maintain external storage solution optimized for use with ProLiant servers. This technology brief describes the MSA500 G2 system and explains how it uses proven industry-standard technology to provide affordable high availability. This document also describes technologies used in the MSA500 G2 system. After reading this technology brief, customers will understand why and where to use an MSA500 G2 system.

## Introduction

In today's Internet economy, more and more companies are using an e-commerce business model. Companies now allow customers to shop, users to perform self-service activities, employees to work, and IT administrators to manage systems around the clock: 24-hour availability is now the norm. As a result, systems availability is critical.

Many companies use clustering solutions to provide high availability. Clustering typically consists of two or more servers connected to common components, such as storage, and provides high availability in a couple of ways. First, work can be distributed across multiple servers, either so that the most capable server can complete it or so that many servers can work on it at once. This allows the system to handle subsequent requests more quickly. Second, one or more servers can be designated to act as backup in case another server fails. Use of a backup server drastically reduces downtime due to failures, and it allows administrators to service the systems without downtime.

The HP Modular Smart Array 500 G2 (MSA500 G2) system is an affordable high availability storage solution for small clusters of HP ProLiant servers. Based on well-known industry-standard technology, the MSA500 G2 is a direct-attached-storage (DAS) solution that is easy to deploy and easy to manage.

---

**Figure 1.** HP MSA500 G2 system



---

The MSA500 G2 system is especially beneficial for environments with remote and distributed workgroups and departments (for example, bank branches and retail stores). It is ideal for environments that do not have an existing Fibre Channel infrastructure, and it is a natural step for customers deploying servers with internal Smart Array controllers.

## Technology overview

The MSA500 G2 is a storage system for 2-, 3-, and 4-node clustering. Its 4U enclosure houses up to 14 Universal hot-pluggable Ultra320, Ultra3, or Ultra2 Small Computer Systems Interface (SCSI) hard disk drives.<sup>1</sup> The MSA500 G2 provides a maximum of 4 terabytes of storage capacity using 300-gigabyte Ultra320 drives. The MSA500 G2 system includes the following standard components:

- One MSA500 G2 controller with 256-MB battery-backed cache, (upgradeable to 512 megabytes)
- One 2-port Ultra320 SCSI input/output (I/O) module (4-port I/O module optional)
- Redundant, hot-pluggable power supplies with fans
- Two power cables
- Two host bus adapters
- Two SCSI cables
- One crossover Ethernet cable

### MSA500 G2 controller

The MSA500 G2 controller (Figure 2) is a hot-pluggable and intelligent Redundant Array of Independent Disks (RAID) controller with integrated status indicators and a 256-megabyte read/write battery-backed cache expandable to 512 megabytes (per controller). One controller ships standard with the HP MSA500 G2 system; a redundant controller is optional. The MSA500 G2 controller is built upon a 32-bit superscalar, reduced instruction set computing (RISC) processor and features a high performance RAID coprocessor designed by HP to enhance performance of RAID parity write operations for RAID 5 and Advanced Data Guarding (ADG).

**Figure 2.** Hot-pluggable MSA500 G2 controller



Featuring SCSI Ultra320 technology that offers twice the speed of previous generation protocols, the MSA500 G2 dramatically reduces the amount of time required for a drive rebuild. Reducing the critical period of drive rebuild time lowers the chance of losing data from a multiple drive failure.

The MSA500 G2 controller features battery-backed write cache to protect posted-write data in the event that external power to the system is interrupted, a host computer fails, or the controller itself fails. The write cache can protect data for up to three days. In a system with a single MSA500 G2 controller, if the controller fails, the cache module can be removed and installed on a replacement controller (within three days with no loss of data), and the posted-write data will be flushed. The write cache can protect data for up to three days. Redundant, maintenance-free nickel metal hydride batteries are used on the battery-backed cache to improve reliability.

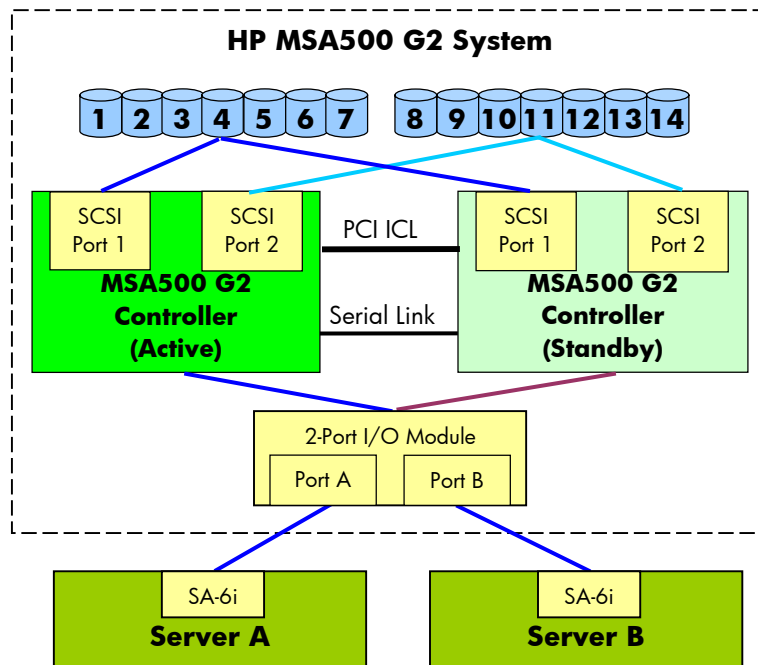
<sup>1</sup> For more information or for details on supported configurations, go to [www.hp.com/products/sharedstorage](http://www.hp.com/products/sharedstorage) or refer to the Quickspecs for the Smart Array Cluster Storage system.

### Redundant controller operation

When redundant controllers are used in the MSA500 G2 solution, the controllers operate in an active/standby configuration, meaning that only one controller at a time can accept commands. However, the write cache of the standby controller is continuously synchronized with the posted-write cache of the primary controller through a high-speed, low-latency Peripheral Component Interconnect (PCI) bus inter-controller link (ICL). Write transactions are not considered complete until both caches are updated. Should the active controller fail, the standby controller with its up-to-date cache will assume control without loss of data or downtime.

To provide better availability, the two controllers communicate through a separate serial link (as shown in Figure 3) so that a failed controller can be distinguished from a failed PCI ICL.

**Figure 3.** Smart Array storage cluster architecture with HP MSA500 G2



Since the MSA500 G2 system is based on SCSI protocols, customers can create high availability solutions without complex network infrastructures. The MSA500 G2 system can attach to Smart Array 642 (SA-642) or embedded Smart Array 6i (SA-6i) controllers so additional connectivity investments are not required for deployment. The SCSI-based storage system eliminates the need for a costly Fibre Channel infrastructure of hubs, switches, host bus adapters (HBAs), and cables. Moreover, SCSI is a familiar, widely adopted protocol; therefore, deploying the MSA500 G2 system does not require the presence of specialized storage professionals or additional expensive specialized training for IT personnel.

In clusters with non-intelligent enclosures or “just a bunch of drives” (JBOD), the RAID controllers are located in each of the clustered servers. Unlike the MSA500 G2 system design, these controllers are not synchronized and have no mechanism to keep cache data coherent. Unsynchronized controllers cannot implement write-caching, the lack of which seriously degrades the performance of the JBOD cluster.

## Recovery ROM and ROM cloning

Each MSA500 G2 controller contains a read-only memory (ROM) that contains its operating firmware. The Recovery ROM feature stores two complete firmware images in the ROM: one active image and one backup image. During the power-up sequence, the controller checks both firmware images to ensure they are valid. The last-known valid image will be used to operate the controller, and the controller will copy this image to replace any invalid firmware image. This auto-flashing ensures that the controller always has a valid firmware image. The controller performs this auto-flashing automatically with no need for user intervention.

In a redundant controller configuration, both MSA500 G2 controllers must execute the same firmware version. During the power-up sequence, the firmware versions on both controllers are compared. If they are not the same, the ROM cloning function will prompt the user (via the LCD panel) to allow cloning to occur.

---

### NOTE:

If the user does not respond to the ROM cloning prompt within 60 seconds, ROM cloning will not occur and the controller with the older firmware will be disabled.

---

Once the copy is complete, the modified controller will be automatically reset and can then begin redundant operation. The controller performs this functionality automatically and without user intervention.

## I/O module

The 2-port Ultra320 SCSI I/O module that comes standard on the MSA500 G2 system includes two very-high-density SCSI connectors (port A, port B) for connection to the servers and also includes an integrated environmental monitoring unit (EMU). The EMU performs several functions, including monitoring the status of the power supplies and fans within the enclosure, detecting cabling errors, and monitoring disk drive removal and insertion. A 4-port I/O module is available as a part of a high availability option kit (described later).

## Power supplies and cooling

The MSA500 G2 system includes dual hot-pluggable power supplies with fans. The dual power supplies provide redundancy and eliminate a single point of failure. In the unlikely event of a fan or power supply failure, one of the power supply/fan units can be replaced without bringing the cluster down. The Universal power supplies include an auto-sensing capability that allows them to detect 100-volt to 240-volt and 40-Hertz to 60-Hertz power.

## Host bus adapters

The MSA500 G2 uses Smart Array 6i (SA-6i) or Smart Array 642 (SA-642) controllers as host bus adapters (HBAs) to connect to the host servers. The SA-6i and SA-642 controllers not only act as HBAs but also manage the failover function of the external Smart Array Cluster Storage controllers and eliminate the requirement for a slower OS-based utility to perform this function.

The SA-6i or SA-642 controller initiates all failovers. When used with firmware version 1.78 or higher, the HBA functionality of these controllers provides the intelligence to recognize that a MSA500 G2 controller has failed when a command does not complete. The SA-6i or SA-642 controller then issues a command that instructs the redundant MSA500 G2 controller to become active. Refer to [www.hp.com/products/sharedstorage](http://www.hp.com/products/sharedstorage) or the server Quickspecs for more information on the specific controllers supported by each server.


## Configuration flexibility

The MSA500 G2 system is a flexible shared storage system for ProLiant servers. It offers deployment versatility to meet customer storage needs for distributed applications and remote sites. It is specifically designed for quick-and-easy deployment with ProLiant servers. In the standard shipping configuration that includes a 2-port I/O module and two host adapters, the MSA500 G2 supports single connections with two hosts configured as either independent or clustered servers.

### MSA500 G2 High Availability Kit

The MSA500 G2 High Availability Kit allows the MSA500 G2 to be deployed as direct-attached storage (DAS) for four server nodes. The nodes may be clustered, independent, or mixed, making the MSA500 G2 ideal for density-optimized ProLiant servers with a limited number of internal disk drive slots that need greater storage capacities. The kit (Figure 4) includes a 4-port I/O module to replace the standard 2-port I/O module, an MSA500 G2 controller for redundant operation, two host adapters, multipath software, and SCSI cables (not shown).

**Figure 4.** HP MSA500 G2 High Availability Kit Components




The high availability kit offers the following configuration options:

- 2-node multipath – Customers can enhance fault tolerance by adding dual paths from two host servers to the shared storage.
- 4-node shared storage – By connecting four independent servers to the MSA500 G2 system, customers can take advantage of the Selective Storage Presentation (SSP) capability.
- 4-node cluster – Clustered servers for Linux® operating systems may be attached to the MSA500 G2 system to meet high availability requirements.
- 4-node cluster/shared storage mix – The MSA500 G2 provides the flexibility of providing storage for a combination of a single 2-node cluster and two independent servers.

## 2-node multipath configuration

Customers can enhance fault tolerance and I/O performance by implementing dual paths with two host servers (as shown in Figure 5). Multipaths provide an additional level of fault tolerance by adding redundancy to both the HBAs in the host servers and the SCSI cables connecting the host servers to the shared storage. Multipaths also allow for configuring static load balancing between the HBAs to improve I/O performance.


**Figure 5.** 2-Node multipath configuration



## 4-node shared storage configuration

Using the high availability kit and a 4-node shared storage configuration allows customers to directly attach up to four independent servers to the HP MSA500 G2 system (Figure 6). Using the Selective Storage Presentation (SSP) utility that ships standard with the HP MSA500 G2 system, customers can assign a specific amount of storage capacity to each of the servers. Implementing SSP and sharing storage among multiple host servers allows for maximum capacity utilization. In addition, online volume extension and online capacity expansion capabilities enable users to add capacity without downtime.

**Figure 6.** MSA500 G2 supporting four independent servers






#### 4-node mixed configuration

The MSA500 G2 with the optional 4-port I/O module allows the flexibility to support hosts that are clustered, independent, or a mix of both. In Figure 7, Hosts A1 and A2 are clustered through an Ethernet crossover cable connection, while Hosts B1 and B2 operate as independent servers. Alternate configurations could have three or all four hosts clustered.

**Figure 7.** 4-node mixed configuration



#### ProLiant DL380 G4 Packaged Cluster system

HP offers a pre-packaged cluster system that includes the MSA500 G2 and two ProLiant DL380 G4 servers in a combined 8U fixture (Figure 8) that greatly simplifies installation and deployment. The ProLiant DL380 G4 servers included in the packaged cluster system ship standard with HP System Insight Manager as well as HP integrated Lights-Out technology for full local or remote management capability.

**Figure 8.** ProLiant DL380 G4 packaged cluster with MSA500 G2 system



The packaged cluster system supports the following cluster software:

- Microsoft® Cluster Service
- Novell Cluster Service
- HP Service Guard Cluster Service for Linux
- SteelEye's Lifekeeper for Linux

For more information on the HP ProLiant DL380 G4 Packaged Cluster system go to the following URL:  
<http://h18004.www1.hp.com/solutions/enterprise/highavailability/dl380/index.html>

## Selective Storage Presentation

Selective Storage Presentation (SSP) is controller firmware that enables administrators to control whether host servers running multiple applications on the MSA500 G2 can access data stored on the MSA500 G2 system. Administrators select which server host or hosts can access the stored data, restricting access as needed to assure data integrity and security. Access control can be granted down to a logical volume level.

## Availability

The MSA500 G2 system offers increased availability and uptime by means of complete hardware redundancy, optimized RAID protections, highly reliable array controllers, and redundant I/O paths (when operating the optional Smart Array Multipath software). The MSA500 G2 system supports several levels of fault tolerance:

- Advanced Data Guarding (ADG)
- Distributed data guarding (RAID 5)
- Drive mirroring (RAID 1 and 1+0 striped mirroring)
- Global online spares

### Advanced Data Guarding

ADG is the highest level of RAID fault tolerance. Two sets of parity data are written across multiple drives, and simultaneous write operations are allowed. This level of fault tolerance can withstand two simultaneous drive failures without downtime or data loss.

### Distributed data guarding

In distributed data guarding, parity data is written across multiple drives, and simultaneous write operations are allowed. This level of protection can withstand one drive failure without downtime.

### Drive mirroring

When using drive mirroring, half of the drive array is allocated to data and the other half to mirrored data. This level of protection provides the highest level of performance and can withstand multiple drive failures as long as they are not in the same mirrored pair of drives.

### Online spares

One or more online spares can be assigned per array and can be used across multiple arrays. If a drive failure occurs, recovery begins immediately using the online spare drive; data is reconstructed automatically from the remaining drives and parity data.

## Management technologies

Remote management capability allows customers to configure and maintain servers and storage without onsite staff. HP management tools such as HP System Insight Manager and the Array Configuration Utility (ACU) offer superior management capabilities to simplify installation and management. The ACU maps the identifications of server host adapters to connection names and sets up access control lists for logical drives based on the identified adapters. Each logical drive on the controller has an access control list that identifies the server host adapters granted access to the drive. If a server attempts to send commands to a logical drive without access authority, the controller rejects the command.

HP remote management products allow IT administrators to perform all of the “in-front-of-the-server” functions without physically visiting the connected servers and storage systems. For example, administrators can remotely perform tasks such as deployment, configuration, maintenance, rebooting servers, troubleshooting, disaster recovery, and application installation. HP Remote Insight Lights-Out Edition II (RiLOE II) and integrated Lights-Out (iLO) technologies deliver powerful, hardware-based remote administration and control of ProLiant servers from a standard web browser. These tools provide maximum security, availability, and control over business-critical ProLiant servers and storage without the need to physically visit the systems.

All Smart Array products use a standard set of management and utility software. Use of standard tools minimizes the customer’s total cost of ownership by reducing training requirements. The ACU allows host-based configuration and remote configuration, and HP System Insight Manager provides remote monitoring. The management tools enable:

- Online RAID level migration, which enables administrators to change the existing RAID level without destroying data or volume information.
- Online capacity expansion for DAS, which allows administrators to add storage to an operating Smart Array Cluster Storage system without expensive server downtime.
- Online volume extension for DAS, which allows administrators to easily grow capacity of existing, logical volumes.
- Global online spare, which reduces the risk of data loss by facilitating automatic rebuilds after a drive failure.
- Logical drive capacity extension, which enables administrators to increase the size of existing logical drives online under Microsoft® Windows NT® and offline for other operating systems.
- Online stripe size migration, which allows administrators to change the stripe sizing of a logical volume.

In addition, the ACU tools enable administrators to dynamically allocate read/write cache sizes, to select expansion and rebuild priority, and to clone controller firmware.


## Hardware and software compatibility

The MSA500 G2 system is compatible with various types of ProLiant servers and operating systems, depending on whether a direct attached shared storage configuration or a clustered configuration is employed. A complete compatibility matrix for the MSA500 G2 is available at the following URL: <ftp://ftp.compaq.com/pub/products/servers/proliantstorage/sharedstorage/msa500g2-compatibility.pdf>

## Performance enhancement

The MSA500 G2, employing RAID 5 technology and a large write cache, provides a substantial increase in performance over earlier systems. Figure 9 illustrates the increase in I/O throughput provided by the MSA500 G2 over the earlier MSA500 G1 system in performing random reads/writes.

**Figure 9.** I/O throughput performance gain of the MSA500 G2 over the MSA500 G1



System test configuration:

MSA500 G1 w/FW v1.80, MSA500 G2 w/FW v1.14, both systems using Seagate Ultra320 36 GB drives (14), ProLiant DL-380 storage system, two Intel XEON processors @ 2.8 GHz, 512 MB RAM, PCI-X bus (64/100), Windows Enterprise edition

## Data migration

Data stored in ProLiant servers protected by Smart Array PCI controllers can be seamlessly migrated to the MSA500 G2 system. Data stored on Universal disk drives using Smart Array controllers can be easily migrated to the MSA500 G2 system by simply removing the drives and inserting them into the MSA500 G2 system. Existing data and configuration will remain intact. The RAID set and data will be preserved, allowing complete migration in minutes, not hours.

## DAS-to-SAN conversion

The MSA500 G2 system offers DAS-to-SAN (DfS) technology, a quick-and-easy way to convert the system to the Modular Smart Array 1000 (MSA1000) when customer requirements change. To convert the MSA500 G2 system to a SAN, customers simply remove the MSA500 G2 controllers (as shown in Figure 10), replace them with the StorageWorks MSA1000 controllers, and add a Fibre Channel I/O board or Fibre Channel switch board and a Fibre Channel Host Bus Adapter.

---

**NOTE:**

The MSA1000 controller does not support the 4-port I/O module. If a 4-port I/O module is installed in a unit to be upgraded with a MSA1000 controller, the 4-port I/O module must be replaced with the 2-port I/O module that came with MSA500 G2.

---

---

**Figure 10.** Hot-pluggable MSA500 G2 controller



## Conclusion

Using industry-standard protocol and components, the MSA500 G2 system is an easy-to-deploy storage solution that provides high availability and high performance to customers with low initial investment requirements.

Since the MSA500 G2 system is based on SCSI protocols, customers can leverage existing infrastructure to create high availability solutions. Moreover, SCSI is a familiar, widely adopted protocol; therefore, deploying the MSA500 G2 system does not require the presence of highly skilled storage professionals or expensive specialized training for IT personnel.

The modular, scalable design of the MSA500 G2 system provides an extremely flexible platform. Instead of making a large up-front investment, customers can buy what they need today and purchase additional capacity and performance as data storage needs grow. Additional servers, drives, a redundant controller, and more battery-backed cache can be easily added as needs dictate.

The MSA500 G2 system supports Universal drives, so customer investments in SCSI technology are protected and cost is minimized as new drives are added. Sharing a common Universal drive form factor between server and storage can significantly reduce needs for hardware and spare inventory and reduce overall storage management and business costs.

The MSA500 G2 system uses a standard, integrated set of management and utility software. These tools consistently lower the cost of ownership by reducing training and technical expertise necessary to install and maintain storage.

The MSA500 G2 system is an exceptional solution for simple, affordable clustering or shared storage. It demonstrates HP's commitment to developing products that provide performance and value in the enterprise.

## For more information

For more information on HP StorageWorks Storage Array products, access the HP website at: <http://h18006.www1.hp.com/storage/arraysystems.html>

## Call to action

To help us better understand and meet your needs for ISS technology information, please send comments about this paper to: [TechCom@HP.com](mailto:TechCom@HP.com)

© 2004 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft, Windows, and Windows NT are U.S. registered trademarks of Microsoft Corporation.

Linux is a U.S. registered trademark of Linus Torvalds.